

# A novel and scalable Multigrid algorithm for many-core architectures

Julián Becerra-Sagredo · Carlos Málaga ·  
Francisco Mandujano

Received: date / Accepted: date

**Abstract** Multigrid algorithms are among the fastest iterative methods known today for solving large linear and some non-linear systems of equations. Greatly optimized for serial operation, they still have a great potential for parallelism not fully realized. In this work, we present a novel multigrid algorithm designed to work entirely inside many-core architectures like the graphics processing units (GPUs), without memory transfers between the GPU and the central processing unit (CPU), avoiding low bandwidth communications. The algorithm is denoted as the high occupancy multigrid (HOMG) because it makes use of entire grid operations with interpolations and relaxations fused into one task, providing useful work for every thread in the grid. For a given accuracy, its number of operations scale linearly with the total number of nodes in the grid. Perfect scalability is observed for a large number of processors.

**Keywords** Parallel multigrid · GPU · CUDA

**PACS** 65F10 · 65N22 · 65N55 · 65Y05 · 65Y10

## 1 Introduction

The multigrid algorithm [1,2] is one of the fastest methods for solving linear and non-linear systems of equations derived from a variety of problems, like numerical discretizations of partial differential equations and non-linear variational problems [3,4]. The main idea is to accelerate the convergence of an iterative method using a hierarchy of nested discretizations to perform

---

J. Becerra-Sagredo  
Inst. Ingeniería Ed. 5 Cub. 307, UNAM, Ciudad Universitaria, 04510 México D.F., México  
Tel.: +525-562-24855  
E-mail: juliansagredo@gmail.com

C. Málaga · F. Mandujano  
Depto. Física, Fac. Ciencias, UNAM, Ciudad Universitaria, 04510 México D.F., México

resolution-dependent corrections that are passed between levels using interpolations. Multigrid has the main advantage over other methods that when used to solve the problems with a given accuracy, its number of operations often scale linearly with the number of discrete nodes used. It has been widely studied and optimized for sequential operation during the last decades, providing very fast convergent algorithms that mostly lack fine grain parallelism because of the sequential use of updates during the same iteration.

In the last few years, hardware aspects have led to a paradigm change in the design of numerical methods. Performance improvements now come by parallelization and specialization and not any more by frequency scaling. Many-core parallel architectures, like GPUs, seem ideally suited for the acceleration of iterative methods like multigrid, given that its basic operations are essentially local and explicit.

Early works on multigrid for GPUs date back to 2003 [5,6], when GPUs started to outperform CPUs and control over the operations and memories of the GPU were available using programmable vertex shaders. These implementations are many-core maps of a classic multigrid algorithm using recursive V-cycles. Additional interesting works in this direction are found in [7,8,9,10], and more recently using Nvidia's compute unified device architecture (CUDA) in [11,12,13,14,15,16]. Nevertheless, a logical consequence of the use of a classic multigrid in many-core architectures is the appearance of performance penalties for coarse grids where the number of independent operations is reduced and memory latencies cannot be hidden using multithreading.

In order to avoid these penalties, a hybrid approach is preferred by some authors [17,18,19] where the CPU is used for serial matrix inversions over coarse grids and the GPU for computing fine-grid relaxations. This is a natural way to avoid the low parallelism of the coarse grids but with the penalty of the memory exchange between the random access memory (RAM) of the CPU and that of the GPU, which can be several orders of magnitude lower in bandwidth than the GPU fast off-chip memory.

In this work, we modify the classic multigrid algorithm giving priority to keeping the data on the GPU, exploiting its fast memory levels and maintaining a constant occupancy. The result is a simple, fully parallel and perfectly scalable multigrid solver that works entirely inside the GPU without the need to communicate data to the CPU. We prove the perfect scalability of the algorithm using several GPUs with different numbers of processors. The extension to multiple GPUs and clusters is beyond the scope of this work.

## 2 A high occupancy many-core multigrid (HOMG)

Linear multigrid algorithms are iterative methods to solve large linear systems of the form

$$A\mathbf{u} = \mathbf{f}, \quad (1)$$

using a hierarchy of discretizations in nested grids where level-dependent restrictions  $P$ , relaxations  $R$ , and interpolations  $S$  are performed in a variety of cycles.

Classic V-cycles  $V(\eta_1, \eta_2)$  consist of the recursive application of the following two-level correction scheme: (a)  $\eta_1$  pre-smoothing relaxations  $v_h \leftarrow R(\eta_1)v_h$ ; (b) the restriction of the residual  $r_h = f_h - A_h v_h$  to the next coarser grid  $r_{2h} \leftarrow P r_h$ ; (c) the solution of the residual equation  $A_{2h} e_{2h} = r_{2h}$ ; (d) the interpolation of the error to the fine grid and the correction of the solution  $v_h \leftarrow v_h + S e_{2h}$ ; and (e)  $\eta_2$  post-smoothing relaxations  $v_h \leftarrow R(\eta_2)v_h$ . Usually, interpolations are linear, restrictions are linearly weighted averages of neighboring sources and relaxations are one of the variations of Jacobi or Gauss-Sidel methods [1].


A fundamental issue in many-core architectures like GPUs is the efficient use of its fast parallel fetches, optimized in every warp for contiguous memory locations and a large number of threads. Massive multithreading and the fast memory cache levels must be efficiently used to reduce cache misses and memory latencies. With this in mind, it is logical that some problems appear in multigrid algorithms as the grid is coarsened because the number of threads, proportional to the number of points in the grid, will be quite low, sometimes even lower than the number of processors available, and the processors will have very small amounts of work or will idle. Even colored Gauss-Sidel methods [1] could reduce the number of parallel threads and complicate the memory access patterns. The coarsening of the grid and the consequent reduction of workload, basic properties of multigrid, become less attractive under the optic of GPUs, where operations over the entire grid are optimal and inexpensive in one, two and three-dimensional structured grids.

A multigrid algorithm that uses the entire grid for every iteration looks contradictory at first sight because it has more operations than the classic version, but it is attractive in many-core architectures because it keeps a constant optimal memory model fully exploiting its parallel fetching and processing capabilities. One way to provide useful work for every thread during a coarse grid correction, is to fuse relaxations and interpolations in one task routine. Furthermore, we don't need to allocate different meshes because we could always act simultaneously in complementary subsets of the entire grid. One subset of points computes relaxations while the rest perform interpolations of the previous iteration variations. This strategy provides a constant high occupancy of the cores.

The damped Jacobi relaxation [1] is a natural choice for a fully parallel algorithm that only uses data produced in a previous iteration. This method is equivalent to the explicit integration of a discretized version of

$$\frac{\partial u}{\partial t} = \kappa (\mathcal{L}(u) - f), \quad (2)$$

where  $\mathcal{L}(u)$  is any elliptic partial differential operator. For example, if  $\mathcal{L}(u) = \nabla^2 u$  and the Equation (2) is discretized using centered spatial differences,



**Fig. 1** Modified full multigrid cycle (MFMG). The number of iterations per level are doubling for every finer level until reaching a maximum number ( $MaxI$ ).

forward Euler integration in time and nested meshes doubling the spatial discretization, the relaxation at level  $m$  is given by

$$u_{i,j}^{n+1} = u_{i,j}^n + \omega \left( (4u_{i,j}^n - u_{i-l,j}^n - u_{i+l,j}^n - u_{i,j-l}^n - u_{i,j+l}^n) - h_l^2 P_m(f_{i,j}) \right), \quad (3)$$

where  $l = 2^m$ ,  $h_l = 2^m h$  and  $P_m$  is the weighted average restriction to level  $m$ .

The high occupancy multigrid (HOMG) algorithm we propose is quite simple: (a)  $\eta_1$  relaxations-interpolations on the coarsest level  $v_{2^M h} \leftarrow R(\eta_1)v_{2^M h}$  and  $v_h \leftarrow S v_{2^M h}$ ; (b)  $\eta_2$  relaxations-interpolations on the next finer level  $v_{2^{(M-1)} h} \leftarrow R(\eta_2)v_{2^{(M-1)} h}$  and  $v_h \leftarrow S v_{2^{(M-1)} h}$ ; (c) go back to a coarser level or descend to a finer level in a chosen pattern until reaching the finest level.


We have explored many sequences of coarse to fine grid levels and found the most efficient to be the modified full multigrid cycle (MFMG) presented in Fig. 1, with only two iterations in the coarsest level, increasingly doubling them for every finer level until reaching a chosen bound  $MaxI$ .

The restrictions are independent of the cycle and can be done in a pre-processing step using several independent arrays in memory. If memory savings are desirable, the restrictions can be done inside the cycle with just one extra memory array. The restrictions are performed inside the GPU using linear weighted averages organized as a series of nested partial reductions over the entire mesh.

After an MFMG cycle is complete, the residual equation is used to restart the cycles as many times as needed to reach any desired accuracy.

### 3 Power of two size grids

Inherently to their architecture, GPUs perform optimally in grids with a power of two number of points per side while multigrid algorithms are usually limited to nested grids. We have designed a geometric strategy to allow the use of the HOMG in a two-dimensional grid with a power of two points per side. The procedure is illustrated in Fig. 2.



**Fig. 2** Geometric strategy for multigrid in power of two grids. The grid is divided into four subdomains with a power of two number of points per side. Each set of symbols represents the location of the coarse  $2h$  grid in each subdomain. Some of them lie outside the subdomain. During a  $2h$  relaxation-interpolation iteration, relaxations are performed over all the points with a symbol and interpolations on the rest. The interpolations avoid cross-domain data.

The grid is divided into four subdomains with a power of two nodes per side. The coarse grid points are marked for each subdomain relative to the corresponding external corner. These marked points would form a unique set in a nested configuration but here they form four different sets. The coarse subsets need to include central points, if not, some coarse modes, like the zero mode, can not be resolved. Due to the absence of central points we mark the coarse grid points that lie just outside the corresponding subdomain. We perform relaxations in all the marked points and interpolations on the rest of the fine grid. The interpolations make use only of points in the same subdomain to avoid cross-domain interpolations.

## 4 Results

We test our algorithm solving the two-dimensional Poisson equation

$$\nabla^2 u = f \quad (4)$$

on the unit square  $\Omega = [0, 1] \times [0, 1]$  with Dirichlet boundary condition  $u = 0$  on the edge  $\partial\Omega$ .

We choose the polynomial analytic solution


$$u(x_1, x_2) = -x_1^2 x_2^2 (1 - x_1^2)(1 - x_2^2) \quad (5)$$

corresponding to the source

$$f(x_1, x_2) = -2(x_2^2(1 - 6x_1^2)(1 - x_2^2) + x_1^2(1 - 6x_2^2)(1 - x_1^2)). \quad (6)$$

This solution has a wide spectrum with components in all the modes of the grid. This characteristic is quite useful to study the convergence of multigrid algorithms like the HOMG, and has been previously used in [1].

Fig. 3 shows the sequential approximation of the solution with the HOMG algorithm in one MFMG cycle with a maximum of two iterations per level



**Fig. 3** High occupancy multigrid (HOMG) approximation to the solution of Eq. (4) with the polynomial source (6) in a  $128^2$  grid. The approximation is shown for different iterations of the modified full multigrid cycle (MFMG) with a maximum of two iterations per level ( $MaxI = 2$ ).


( $MaxI = 2$ ) and a  $128^2$  grid. During the first iteration we see a spike because only the middle points perform a relaxation and the interpolations have null sources. Remember that the interpolations are delayed one iteration because they are performed in parallel together with the relaxations. The second iteration shows that the relaxations have stalled while the interpolations form the zero mode approximation with pyramid shape. The third iteration shows clearly the next finer level relaxations-interpolations and so on. The MFMG cycle has a total of 31 iterations for a  $128^2$  grid.

The error of the approximations is computed using the normalized  $L_1$ -error


$$Error = \frac{1}{\int_{\Omega} |u| d\Omega} \int_{\Omega} |u - v_h| d\Omega. \quad (7)$$

Fig. 4 shows the convergence of the  $L_1$ -error for the HOMG method for one MFMG cycle and two different choices of the maximum iterations per cycle  $MaxI$ . The results show all the desirable properties of a multigrid algorithm. Cycles with the same  $MaxI$  descend over the same convergence path for different resolutions. Higher resolutions continue the smoothing of the solution obtained with coarser grids. Increasing  $MaxI$  allows the cycles to reach lower errors with the penalty of a large number of iterations per cycle.  $MaxI = 4$  reaches an error of 5% and  $MaxI = 32$  reaches an error of 1.6%. Operations over the entire grid are denoted as a working unit (WU). Given a desired accuracy, the algorithm performs a number of operations linearly proportional to the points in the grid. Remember that after an MFMG cycle is complete, the residual equation is used to restart the cycles as many times as needed to reach any desired accuracy.

We explore the scalability of the HOMG algorithm on GPUs with different architectures and number of processors. The test consists of one MFMG cycle with  $MaxI = 32$  and different resolutions. The GPUs used are all GeForce from Nvidia. We use C language for CUDA as the application programming



**Fig. 4**  $L_1$ -error (Error) convergence of the modified multigrid cycle for different domain sizes and maximum number of iterations per level ( $MaxI$ ). The working units (WU) correspond to entire grid operations.



**Fig. 5** Computing time in seconds for  $MaxI = 32$  in GPUs with different architectures and number of processors (Nprocs). The perfect linear scalability curve is shown for comparison.

interface. Five GPUs of the modified 8-series architecture: a 310M with 16 processors, a 9500GT with 32 processors, a 330M with 48 processors, a 9800GT with 112 processors and a GTX260 with 216 processors. And two GPUs of the GeForce Fermi architecture: a GTX460 with 336 processors and a GTX480 with 480 processors. All the computations are done in single precision.

The scalability results are shown in Fig. 5. The slopes are moderate for a low number of processors and small grids. For a large number of processors and large grids the algorithm reaches perfect scalability.

## 5 Conclusion

In this article, we present a novel multigrid algorithm specially designed to work entirely inside many-core architectures like GPUs without memory transfers between the GPU and the CPU. The algorithm makes use of entire grid operations even for coarse grid corrections. Interpolations and relaxations are fused into one task giving useful work for every thread in the grid. In this way the algorithm has full multithreading and fix memory patterns, allowing the full exploitation of the fast memory models of the GPU, efficiently hiding cash misses and memory latencies.

The algorithm is denoted as the high occupancy multigrid (HOMG) algorithm because multithreading and useful work per thread are kept constantly high. The algorithm is combined with a modified full multigrid cycle (MFMG) to reach a high efficiency. For a given accuracy, the operations of the HOMG scale linearly with the total number of nodes. Perfect scalability is observed for a large number of processors and large grids.

## References

1. Briggs, W.L., Henson, V.E., McCormick, S.F., A Multigrid Tutorial, 2nd edn. Society for Industrial and Applied Mathematics, Philadelphia (2000)
2. Trottenberg, U., Oosterlee, C.W., Schüller, A., Multigrid. Academic Press, San Diego (2001)
3. Hackbusch, W., Multigrid methods and applications. Springer, Heidelberg (1985)
4. McCormick, S.F., Multigrid methods, Frontiers in applied mathematics, 2nd edn. Society for Industrial and Applied Mathematics, Philadelphia (1994)
5. Bolz, J., Farmer, I., Grinspun, E., Schroeder, P., Sparse Matrix Solvers on the GPU: Conjugate Gradients and Multigrid, ACM Transactions on Graphics, Vol. 22, pp. 917-924 (2003)
6. Goodnight, N., Woolley, C., Lewin, G., Luebke, D., Humphreys, G., A Multigrid Solver for Boundary Value Problems using Programmable Graphics Hardware, Siggraph 2003. In: Proceedings of SIGGRAPH, pp. 1021-111 (2003)
7. Strzodka, R., Droske, M., Rumpf, M., Image Registration by a Regularized Gradient Flow: A Streaming Implementation in DX9 Graphics Hardware, Computing, Vol. 73, No. 4, pp. 373-389 (2004)
8. Grossauer, H., Thoman, P., GPU-Based Multigrid: Real-Time Performance in High Resolution Nonlinear Image Processing, Computer Vision Systems, Lecture Notes in Computer Science, Vol. 5008/2008, pp. 141-150 (2008)
9. Grady, L., A Lattice-Preserving Multigrid Method for Solving the Inhomogeneous Poisson Equations Used in Image Analysis, Computer Vision, Lecture Notes in Computer Science, Vol. 5303/2008, pp. 252-264 (2008)
10. Sangild, T., Mosegaard, J., An Introduction to GPU Accelerated Surgical Simulation, ISBMS 2006, Lecture Notes in Computer Science, Vol. 4072, pp. 93-104 (2006)
11. Kazhdan, M., Hoppe, H., Streaming Multigrid for Gradient- Domain Operations on Large Images, ACM Trans. Graphics, Vol. 27, No. 3, pp. 1-10 (2008)
12. Feng, Z., Li, P., Multigrid on GPU: Tackling Power Grid Analysis on Parallel SIMT Platforms, Proc. IEEE/ACM Intl Conf., Computer-Aided Design ICCAD 08, pp. 647-654 (2008)
13. Elsen, E., LeGresley, P., Darve, E., Large Calculation of the Flow over a Hypersonic Vehicle Using a GPU, J. Computational Physics, Vol. 227, No. 24, pp. 10148-10161 (2008)
14. Göddeke, D., Strzodka, R., Mohd-Yusof, J., McCormick, P., Buijssen, S., Grajewski, M., Turek, S., Exploring weak scalability for FEM calculations on a GPU-enhanced cluster, Parallel Computing, Vol. 33:10-11. pp. 685-699 (2007)



- 
15. Rehman, T., Haber, E., Pryor, G., Melonakos, J., Tannenbaum, A., 3D nonrigid registration via optimal mass transport on the GPU. *Medical Image Analysis*, Vol. 13, pp. 931-940 (2009)
  16. Haase, G., Liebmann M., Douglas, C., Plank, G., A Parallel Algebraic Multigrid Solver on Graphics Processing Units, *High Performance Computing and Applications*, Lecture Notes in Computer Science, Vol. 5938/2010, pp. 38-47 (2010)
  17. Thoman, P., GPGPU-based Multigrid Methods, Master Thesis. University of Innsbruck (2007)
  18. Göddeke, D., Strzodka, R., Mohd-Yusof, J., McCormick, P., Wobker, H., Becker, C., Turek, S., Using GPUs to Improve Multigrid Solver Performance on a Cluster, *International Journal of Computational Science and Engineering (IJCSE)*, Vol. 4, No. 1, pp. 36-55, (2008)
  19. Göddeke, D., Strzodka, R., Cyclic Reduction Tridiagonal Solvers on GPUs Applied to Mixed Precision Multigrid , *IEEE Transactions on Parallel and Distributed Systems*, Special Issue: High Performance Computing with Accelerators, Vol. 22, No. 1, pp. 22-32 (2010)